WHAT WE CLAIM IS:

1. A method of training acoustic models for use in phonetically spelled word models comprising:

-using a training pronunciation guesser to generate a phonetic spelling, each including a sequence of phonemes, from the text spelling of each of a set of acoustic training words;

-mapping sequences of sound associated with utterances of each of the training words against the corresponding sequence of phonemes defined by the phonetic spelling associated with the training word by the pronunciation guesser; and

-for each of a plurality of said phonemes, using the sounds mapped against a given phoneme in one or more of said phonetic spellings to develop at least one acoustic phoneme model for the given phoneme.

2. A method as in claim 1 further including using the acoustic phoneme models in speech recognition performed against acoustic word models, where such a word model of a given word is composed of a sequence of the acoustic phoneme models corresponding to a phonetic spelling generated for the given word by a recognition pronunciation guesser.

3. A method as in claim 1 wherein 5% or more of the occurrences of vowel phonemes placed in the phonetic spellings of the acoustic training words by the training pronunciation guesser are phonetic spelling errors.

4. A method as in claim 3:

   -further including using the acoustic phoneme models in speech recognition performed against acoustic word models of words, where the acoustic word model of a given word is composed of a sequence of the acoustic phoneme models corresponding to a phonetic spelling generated for the given word by a recognition pronunciation guesser; and

   -wherein the recognition pronunciation guesser would make 50% or more of the same phonetic spelling errors made by the training pronunciation guesser in the acoustic training words if it were to generate phonetic spellings for the set of acoustic training words.

5. A method as in claim 4 wherein the recognition and acoustic training pronunciation guessers are the same pronunciation guesser.

6. A method as in claim 4 wherein the words whose guessed phonetic spellings are used in the speech recognition are peoples' names.

7. A method as in claim 6 wherein the speech recognition is used in telephone name dialing in which the speech recognition of a name is used to select a telephone number associated with that name that can be automatically dialed.

8. A method as in claim 7 wherein the speech recognition and name dialing are performed on a cellphone.

9. A method as in claim 8 further including:

-storing on said cellphone, for each of a plurality of
commands words used to control the cellphone, a
phonetic spelling of the command that comes from a
source more accurate than the recognition
pronunciation guesser; and   ·

-performing speech recognition on a given utterance by
matching it against acoustic word models, each
composed of a sequence of said acoustic phoneme models
corresponding to one of said stored phonetic spellings
of a command word; and

-responding to an indication by the speech recognition
that the given utterance corresponds to the phonetic
spelling of a given one of the command words by
causing the cellphone to perform the given command.

10. A method as in claim 8 further including:

-responding to the entry of a name by a user by having
the recognition pronunciation guesser generate a
phonetic spelling for the user-entered name; and

-using the phonetic spelling of the user-entered name
in the speech recognition.

11. A method as in claim 10 further including:

-for each of a plurality of common names, testing if
the phonetic spelling produced for the name by the
recognition pronunciation guesser is correct; and

-for each of a plurality of said common names which
are found not to have correct phonetic spellings
generated for them by the recognition pronunciation
guesser, storing on said cellphone a phonetic spelling

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 81

of the name that comes from a source more accurate
than the recognition pronunciation guesser;

-wherein said responding to the entry of a name by a
user includes:

    --checking to see if the name is one for which a
    phonetic spelling from the more accurate source
    has been stored;

    --if so, using the more accurate spelling as the
    phonetic spelling for the user entered word in
    speech recognition; and

    --if not, using the recognition pronunciation
    guesser to generate the phonetic spelling of the
    word and using that generated spelling in speech
    recognition.

12. A method as in claim 3 further including training the
training pronunciation guesser by:

    -obtaining the following data for each of a plurality
    of said pronunciation-guesser training words:

        --a textual spelling for the word, comprised of a
        sequence of letters;

        --a relatively reliable phonetic spelling for the
        word, comprised of a sequence of phonemes; and

        --a measure of the frequency with which the word
        occurs; and

    -using the data obtained for each of said
    pronunciation-guesser training words to train the
    pronunciation guesser, including:

        --for each pronunciation-guesser training word,
        mapping the sequence of letters of the training
        word's textual spelling against the sequence of

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 82

phonemes of the relatively reliable phonetic
spelling; and

--using the resulting letter-to-phoneme mappings
to train the pronunciation guesser;

-wherein the using of said letter-to-phoneme mappings
includes varying the weight given to a given letter-
to-phoneme mapping in the training of the
pronunciation guesser as a function of the frequency
measure of the word in which such a mapping occurs.

13. A method as in claim 12 wherein the ratio of the weight
given to a letter-to-phoneme mapping relative to the
frequency of the given word in which the mapping occurs
decreases as the frequency of the given word increases.

14. A method as in claim 1 wherein a majority of said
acoustic phoneme models are multiphone models, each of
which represents the sound of a given phoneme when it
occurs in a given phonetic spelling context defined by one
or more phonemes occurring before or after the given
phoneme in a phonetic spelling.

15. A method as in claim 1 wherein a majority of said
acoustic phoneme models are monophone models in which a
given acoustic model represents the sounds of a given
phoneme in all the phonetic spelling contexts in which it
can occur in said phonetic spellings.

16. A method as in claim 1 wherein the acoustic training
words are English words.

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 83

17. A method as in claim 1 wherein the pronunciation guesser is trained on a representative distribution of names from US phone books.

18. A method as in claim 17 wherein the training pronunciation guesser is sufficiently errorful that 5% or more of the occurrences of vowel phonemes the training pronunciation guesser would placed in the phonetic spellings of such a set of names, if generating their phonetic spellings, would be phonetic spelling errors.

19. A method of making a speech recognition enabled computing system comprising:
    -training a set of acoustic phoneme models by:
        --using a training pronunciation guesser to generate a phonetic spelling, each including a sequence of phonemes, from the text spelling of each of a set of acoustic training words;
        --mapping sequences of sound associated with one or more utterances of each of the training words against the sequence of phonemes defined by the phonetic spelling associated with the training word by the pronunciation guesser; and
        --for each of a plurality of said phonemes, using the sounds mapped against a given phoneme in one or more of said phonetic spellings to develop at least one acoustic phoneme model for the given phoneme; and
    -storing in machine readable memory of the computing system being made the following:
        --recognition pronunciation guessing programming for generating a phonetic spelling, comprised of

a sequence of phonemes, from a textual spelling
of a word;

--said set of acoustic phoneme models, including
at least one for modeling the speech sounds
associated with each phoneme used in the phonetic
spellings generated by the recognition
pronunciation guessing programming;

--speech recognition programming for recognizing
an utterance by scoring the match between a
sequence of the utterance's speech sounds and a
sequence of said acoustic phoneme models
associated with the phonetic spelling of each of
a plurality of words; and

--programming for enabling the speech recognition
programming to perform recognition against a
sequence of said acoustic phoneme models
associated with a phonetic spelling generated by
the pronunciation guessing programming;

-wherein:

--5% or more of the occurrences of vowel phonemes
placed in the phonetic spellings of the acoustic
training words by the training pronunciation
guesser are phonetic spelling errors; and

--the recognition pronunciation guessing
programming would make 50% or more of the same
phonetic spelling errors as are made by the
training pronunciation guesser when generating
phonetic spellings for the acoustic training
words.

20. A method as in claim 19 further including storing in
said machine readable memory programming for:

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 85

-enabling a user to enter the text spelling of a name
into the system in association with an item upon which
the system can perform a given function;

-responding to such a user's entry of a name into the
system by causing the pronunciation guessing
programming to generate a phonetic spelling from the
text spelling of the entered name;

-responding to a user's utterance by having the speech
recognition programming score the match between the
sound of the utterance and sequences of said acoustic
phoneme models corresponding to the phonetic spellings
generated by the pronunciation guessing programming
for each of one or more user entered names; and

-determining whether to perform the given function on
the item associated with a given user-entered name as
a function of the score produced by the speech
recognition programming for the utterance against the
phonetic spelling of the given user-entered name.

21. A method as in claim 20 wherein:

-the item associated with a user-entered name includes
a phone number; and

-the given function is the dialing of the phone number
associated with a user-entered name selected as a
function of the score produced by the speech
recognition programming.

22. A method as in claim 21 wherein the system is a
cellphone.

23. A method as in claim 20:

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 86

-further including storing in said machine readable memory correct phonetic spellings for a plurality of names the recognition pronunciation guessing programming phonetically misspells; and

-wherein said programming for responding to a user's entry of a name into the system includes programming for responding to the user's entry of a given name for which a correct phonetic spelling has been stored by causing said correct phonetic spelling to be used as the phonetic spelling for the given user-entered name in matching performed by the speech recognition programming instead of a phonetic spelling generated for the given name by said recognition pronunciation guessing programming.

24. A method as in claim 23 wherein said speech recognition programming uses the same acoustic phoneme models for a given phoneme in a given phonetic context in said correct phonetic spellings as it uses for the same phoneme in the same phonetic context in phonetic spellings generated by the pronunciation guessing programming.

25. A method as in claim 20 further including storing in said machine readable memory:

-a correct phonetic spelling for each of a plurality of commands;

-command recognition programming for causing the speech recognition programming to perform recognition of utterances against sequences of said acoustic phoneme models corresponding to the stored correct phonetic spellings of said commands; and

-programming for determining whether to perform a
given command as a function of the score produced by
the speech recognition programming of a given
utterance against the correct phonetic spelling of the
given command.

26. A speech recognition system comprising:
   -machine readable memory storing;
      --pronunciation guessing programming for
      generating a phonetic spelling, comprised of a
      sequence of phonemes, from a textual spelling of
      a word;
      --a set of acoustic phoneme models, including at
      least one for modeling the speech sounds
      associated with each phoneme used in the phonetic
      spellings generated by the pronunciation guessing
      programming;
      --speech recognition programming for recognizing
      an utterance by scoring the match between a
      sequence of the utterance's speech sounds and a
      sequence of said acoustic phoneme models
      associated with the phonetic spelling of each of
      a plurality of word models; and
      --programming for enabling the speech recognition
      programming to perform recognition against
      phonetic spellings generated by the pronunciation
      guessing programming;
   -wherein:
      --each of said acoustic models represents a
      phoneme in phonetic context;
      --each of a plurality of said acoustic models is
      a blended acoustic model that represents a given

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 88

phoneme in a given phonetic context as a
distribution of sounds corresponding to
utterances of the given phoneme and utterances of
an associated set of one or more other phonemes;
and

--over the plurality of blended acoustic models,
the relative weight allocated, in a given
acoustic model representing a given phoneme in a
given phonetic context, between sounds of
utterances of the given phoneme and each of the
given phoneme's associated phonemes varies as a
function of the frequency with which the
pronunciation guessing programming places the
given phoneme in a position in a phonetic
spelling in the given phonetic context where the
correct phoneme for the position is,
respectively, the given phoneme and each of said
associated phonemes.

27. A system as in claim 26 wherein said machine readable
memory further stores programming for:

-enabling a user to enter the textual spelling of a
word into the system;

-responding to such a user's entry of a word into the
system by causing the pronunciation guessing
programming to generate a phonetic spelling from the
textual spelling of the entered word; and

-responding to a user's utterance by having the speech
recognition programming score the match between the
sound of the utterance and sequences of acoustic
phoneme models corresponding to the phonetic spellings

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 89

generated by the pronunciation guessing programming for each of one or more user entered words.

28. A system as in claim 27 wherein:
    -said machine readable memory further stores correct phonetic spellings for a plurality of words the pronunciation guessing programming phonetically misspells; and
    -said responding to a user's entry of a word into the system includes responding to the user's entry of a given word for which a correct phonetic spelling has been stored by causing said correct phonetic spelling to be used as the phonetic spelling that is used, in conjunction with said acoustic phoneme models, to represent the given user-entered word in the matching performed by the speech recognition programming instead of a phonetic spelling generated for the given name by said recognition pronunciation guessing programming.

29. A method as in claim 28 wherein said speech recognition programming uses the same blended acoustic phoneme models for a given phoneme in a given phonetic context in said correct phonetic spellings as it uses for the same phoneme in the same phonetic context in phonetic spellings generated by the pronunciation guessing programming.

30. A system as in claim 27 wherein said machine readable memory further stores:
    -a correct phonetic spelling for each of a plurality of commands;

-command recognition programming for causing the
speech recognition programming to perform recognition
of utterances against sequences of said acoustic
phoneme models, including said blended acoustic
phoneme models, corresponding to the stored correct
phonetic spellings of said commands; and

-programming for determining whether to perform a
given command as a function of the score produced by
the speech recognition programming of a given
utterance against the correct phonetic spelling of the
given command.

31. A speech recognition system comprising:

-a pronunciation guesser for generating a phonetic
spelling, comprised of a sequence of phonemes, from a
textual spelling of a word;

-machine readable memory storing a set of acoustic
phoneme models, including at least one for modeling
the speech sounds associated with each phoneme used in
the phonetic spellings generated by the pronunciation
guesser;

-a speech recognizer for recognizing an utterance by
scoring the match between a sequence of the
utterance's speech sounds and a sequence of said
acoustic phoneme models associated with the phonetic
spelling of each of a plurality of word models; and

-circuitry for enabling the speech recognizer to
perform recognition against phonetic spellings
generated by the pronunciation guesser;

-wherein:

--each of said acoustic models represents a
phoneme in a phonetic context;

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 91

--each of a plurality of said acoustic models is a blended acoustic model that represents a given phoneme in a given phonetic context as a distribution of sounds corresponding to utterances of the given phoneme and utterances of an associated set of one or more other phonemes; and

--over the plurality of blended acoustic models, the relative weight allocated, in a given acoustic model representing a given phoneme in a given phonetic context, between sounds of utterances of the given phoneme and each of the given phoneme's associated phonemes varies as a function of the frequency with which the pronunciation guesser places the given phoneme in a position in a phonetic spelling in the given phonetic context where the correct phoneme for the position is, respectively, the given phoneme and each of said associated phonemes.

32. A system as in claim 31 further including circuitry for:

-enabling a user to enter the textual spelling of a word into the system;

-responding to a user's entry of a word into the system by causing the pronunciation guesser to generate a phonetic spelling from the textual spelling of the entered word; and

-responding to a user's utterance by having the speech recognizer score the match between the sound of the utterance and sequences of acoustic models corresponding to the phonetic spellings generated by

the pronunciation guessing programming for each of one
or more user entered words.

33. A system as in claim 32 wherein:
    -said machine readable memory further stores correct
    phonetic spellings for a plurality of words the
    pronunciation guesser phonetically misspell; and
    -said responding to a user's entry of a word into the
    system responds to the user's entry of a given word
    for which a correct phonetic spelling has been stored
    by causing said correct phonetic spelling to be used
    as the phonetic spelling for the given user-entered
    word in the matching performed by the speech
    recognizer.

34. A method as in claim 33 wherein said speech recognizer
uses the same blended acoustic phoneme models for a given
phoneme in a given phonetic context in said correct
phonetic spellings as it uses for the same phoneme in the
same phonetic context in phonetic spellings generated by
the pronunciation guesser.

35. A system as in claim 32:
    -wherein said machine readable memory further stores a
    correct phonetic spelling for each of a plurality of
    commands; and
    -said system further includes:
        --command recognition circuitry for causing the
        speech recognizer to perform recognition of
        utterances against sequences of said acoustic
        phoneme models corresponding to the stored
        correct phonetic spellings of said commands; and

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 93

--circuitry for determining whether to perform a
given command as a function of the score produced
by the speech recognizer for a given utterance
against the correct phonetic spelling of the
given command;

-wherein said speech recognizer uses the same blended
acoustic phoneme models for a given phoneme in a given
phonetic context in said correct command phonetic
spellings as it uses for the same phoneme in the same
phonetic context in phonetic spellings generated by
the pronunciation guesser.

36. A system as in claim 31 wherein:

-the pronunciation guesser is such that it would
produce phonetic spellings in which 5% or more of the
individual occurrences of vowel phonemes are phonetic
misspellings when generating the phonetic spellings of
a given vocabularyVST02-1 for which the pronunciation
guesser has been trained to generated phonetic
spellings;

-each of said acoustic models represents a phoneme in
a phonetic context;

-each of a set of said acoustic models, including at
least one acoustic model for each of a plurality of
the vowel phonemes used by the pronunciation guesser,
is a blended acoustic model that represents a given
phoneme in a given phonetic context as a distribution
of sounds corresponding to utterances of the given
phoneme and utterances of an associated set of one or
more other phonemes; and

-over the plurality of blended acoustic models, the
relative weight allocated, in a given acoustic model

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 94

representing a given phoneme in a given phonetic context, between sounds of utterances of the given phoneme and each of the given phoneme's associated phonemes is correlated with the frequency with which the pronunciation guesser would place, when generating phonetic spelling for the given vocabulary, the given phoneme in a position in a phonetic spelling within the given phonetic context where the correct phoneme for the position is, respectively, the given phoneme and each of said associated phonemes.

37. A speech recognition system comprising:
   -machine readable memory storing;
      --pronunciation guessing programming for generating a phonetic spelling, comprised of a sequence of phonemes, from a textual spelling of a word;
      --a set of acoustic phoneme models, including at least one for modeling the speech sounds associated with each phoneme used in the phonetic spellings generated by the pronunciation guessing programming;
      --speech recognition programming for recognizing an utterance by scoring the match between a sequence of the utterance's speech sounds and a sequence of said acoustic phoneme models associated with the phonetic spelling of each of a plurality of word models; and
      --programming for enabling the speech recognition programming to perform recognition against phonetic spellings generated by the pronunciation guessing programming;

-wherein:

    --the pronunciation guessing programming would
produce phonetic spellings in which 5% or more of
the individual occurrences of vowel phonemes are
phonetic misspellings when generating the
phonetic spellings of a given vocabulary for
which the pronunciation guesser has been trained
to generated phonetic spellings;

    --each of said acoustic models represents a
phoneme in a phonetic context;

    --each of a plurality of said acoustic models,
including at least one acoustic model for at
least a plurality of vowel phonemes used by the
pronunciation guessing programming, is a blended
acoustic model that represents a given phoneme in
a given phonetic context as a distribution of
sounds corresponding to utterances of the given
phoneme and utterances of an associated set of
one or more other phonemes; and

    --over the plurality of blended acoustic models,
the relative weight allocated, in a given
acoustic model representing a given phoneme in a
given phonetic context, between sounds of
utterances of the given phoneme and each of the
given phoneme's associated phonemes is correlated
with the frequency with which the pronunciation
guessing programming would place, when generating
phonetic spelling for the given vocabulary, the
given phoneme in a position in a phonetic
spelling within the given phonetic context where
the correct phoneme for the position is,

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 96

respectively, the given phoneme and each of said associated phonemes.

38. A speech recognition system as in Claim 37 wherein a majority of said blended acoustic models are multiphone models, each of which represents the sound of a given phoneme when it occurs in a given phonetic spelling context defined by one or more phonemes occurring before or after the given phoneme in a phonetic spelling.

39. A speech recognition system as in Claim 37 wherein a majority of said blended acoustic models are non-multiphone models in which a given acoustic model represents the sounds of a given phoneme in all the phonetic spelling contexts in which it can occur in said phonetic spellings.

40. A system as in claim 37 wherein said machine readable memory further stores programming for:
    -enabling a user to enter the text spelling of a name into the system in association with an item upon which the system can perform a given function;
    -responding to such a user's entry of a name into the system by causing the pronunciation guessing programming to generate a phonetic spelling from the text spelling of the entered name;
    -responding to a user's utterance by having the speech recognition programming score the match between the sound of the utterance and sequences of said acoustic phoneme models corresponding to the phonetic spellings generated by the pronunciation guessing programming for each of one or more user entered names; and

-determining whether to perform the given function on the item associated with a given user-entered name as a function of the score produced by the speech recognition programming for the utterance against the given user-entered name.

41. A system as in claim 40 wherein a user-entered name is a person's name.

42. A system as in claim 40 wherein:
-the item associated with a user-entered name includes a phone number; and
-the given function is the dialing of the phone number associated with the user-entered name selected by the speech recognition programming.

43. A system as in claim 42 wherein the system is a cellphone.

44. A system as in claim 40 wherein:
-said machine readable memory further stores correct phonetic spellings for a plurality of names the pronunciation guessing programming phonetically misspell; and
-said responding to a user's entry of a name into the system responds to the user's entry of a given name for which a correct phonetic spelling has been stored by causing said correct phonetic spelling to be used as the phonetic spelling for the given user-entered name in the matching performed by the speech recognition programming.

45. A method as in claim 44 wherein said speech recognition programming uses the same blended acoustic phoneme models for a given phoneme in a given phonetic context in said correct phonetic spellings as it uses for the same phoneme in the same phonetic context in phonetic spellings generated by the pronunciation guessing programming.

46. A system as in claim 40 wherein said machine readable memory further stores:

 -a correct phonetic spelling for each of a plurality of commands;

 -command recognition programming for causing the speech recognition programming to perform recognition of utterances against sequences of said acoustic phoneme models, including said blended acoustic phoneme models, corresponding to the stored correct phonetic spellings of said commands; and

 -programming for determining whether to perform a given command as a function of the score produced by the speech recognition programming of a given utterance against the correct phonetic spelling of the given command.

47. A system as in claim 37 wherein a blended acoustic phoneme model representing a given phoneme in a given phonetic context does so without representing which portions of the model's blended distribution of speech sounds are associated with the given phoneme and which are associated with one or more of the given phoneme's associated phonemes.

48. A system as in claim 37 wherein said machine readable memory further stores:

-a pure acoustic phoneme model associated with each of a plurality of phonemes, each of which represents the sound of a given phoneme in a phonetic context with less blending from other phonemes than a corresponding blended acoustic phoneme model for the phoneme;

-for each of said blended acoustic phoneme models, a representation of the relative blending weights to be given to the model's given phoneme and to each of the given phoneme's associated phonemes in the blended acoustic model; and

-programming for creating, for each given one of a plurality of blended acoustic phoneme models, a representation for use by the speech recognition programming of the blend between the model's given phoneme and the given phoneme's associated phonemes from a combination of the pure acoustic phoneme models corresponding to the given phoneme and the given phoneme's associated phonemes, based on the representation of relative blending weights stored for the given blended acoustic model; and

49. A system as in claim 48 wherein said programming for creating said blended a representation for use by the speech recognition programming of the blended acoustic phoneme model of a given phoneme creates the blended representation of the speech sounds associated with utterances of the given phoneme and the given phoneme's associated phonemes that does not separately represent which portions of the blended distribution of speech sounds

are associated with the given phoneme and which portions
are associated with one or more of the given phoneme's
associated phonemes.

50. A system as in claim 48 wherein said programming for ·
creating said blended a representation for use by the
speech recognition programming of a given blended acoustic
phoneme model of a given phoneme does so by causing the
speech recognition programming to compare the portion of an
utterance that is mapped against the given blended acoustic
phoneme model in a given phonetic spelling against the pure
acoustic phoneme models of the given phoneme and the given
phoneme's associated phonemes.

51. A system as in claim 50 wherein the score of the match
against pure models of the given phoneme and the given
phoneme's associated phonemes is a function not only of the
degree of match against the pure model of such phonemes,
but also of the relative blending weights stored in
association with each of those phonemes.

52. A system as in claim 48 wherein said machine readable
memory further stores programming for responding to one or
more training utterances of words by a user of the system
by:
      --mapping the sounds of said one or more training
      utterances against word models, where each such word
      model includes a correct phonetic spelling and a
      sequence of the one or more pure acoustic phoneme
      models associated with said phonetic spelling;
      --altering each pure acoustic phoneme models against
      which a portion of one or more utterances is mapped to

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. – Please
contact Ed Porter at 671-494-1722 with any questions   Page 101

better represent the training utterance sounds mapping
against the pure acoustic phoneme model; and
--causing the programming for creating the
representation of the blend between a blended acoustic
phoneme model's given phoneme and the given phoneme's
associated phonemes to create such a blended
representation from a combination of pure acoustic
phoneme models that have been altered in response to
said training utterances.

53. A method of training a pronunciation guesser
comprising:
    -obtaining the following data for each of a plurality
    of said pronunciation-guesser training words:
        --a textual spelling for the word, comprised of a
        sequence of letters;
        --a phonetic spelling for the word, comprised of
        a sequence of phonemes; and
        --a measure of the frequency with which the word
        occurs;
    -using the data obtained for each of said
    pronunciation-guesser training words to train the
    pronunciation guesser, including:
        --for each pronunciation-guesser training word,
        mapping the sequence of letters of the training
        word's textual spelling against the sequence of
        phonemes of the phonetic spelling for the
        training; and
        --using the resulting letter-to-phoneme mappings
        to train the pronunciation guesser;

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 102

-wherein the using of said letter-to-phoneme mappings
includes varying the weight given to a given letter-
to-phoneme mapping in the training of the
pronunciation guesser as a function of the frequency
measure of the word in which such a mapping occurs.

54. A method as in Claim 53 wherein said words are names

55. A method as in Claim 53 wherein said pronunciation
guesser being trained is a D-Tree pronunciation guesser.

56. A method as in Claim 53 wherein the ratio of training
weight to frequency is less for more frequent words than
for less frequent words.

57. A method as in Claim 56 wherein the training weight
varies as function of frequency raised to power less than 1

58. A method as in Claim 53 further including:
    -using the pronunciation guesser to generate a
    phonetic spelling of a word; and
    -using a text-to-speech synthesizer to generate the
    sound of the word from said phonetic spelling.

59. A method as in Claim 58 wherein:
    -said using of the pronunciation guesser includes
    responding to the entry of a name by a user by having
    the pronunciation guesser generate a phonetic spelling
    for the user-entered name;
    -said method further includes performing speech
    recognition against acoustic word models of the user-
    entered names, each of which is composed of a sequence

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 103

of acoustic phoneme models corresponding to a phonetic
spelling generated for the name by the pronunciation
guesser, to select a given word as best matching an
utterance; and

-the use of the text-to-speech synthesizer includes
indicating to a user which name has been selected by
the speech recognition by having the text-to-speech
synthesizer generate the sound of the recognized name.

60. A method as in Claim 59 wherein:
    -said user entered names are names associated with
    phone numbers;
    -said method further includes responding to the
    selection of a name by the speech recognition by
    automatically dialing the phone number associated with
    the recognized name; and
    -wherein the indicating to a user of which name has
    been recognized is performed to inform the user of
    which name the speech recognition has selected.

U.S. Non-Provisional Patent application entitled "Training And Using Pronunciation
Guessers for Speech Recognition" filed on 10/10/03 by Laurence S. Gillick et al. - Please
contact Ed Porter at 671-494-1722 with any questions   Page 104